



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
11.03.2015 Bulletin 2015/11

(51) Int Cl.:
G06F 1/32 (2006.01)

(21) Application number: **14174593.5**

(22) Date of filing: **26.06.2014**

(84) Designated Contracting States:
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR**
Designated Extension States:
BA ME

- **Wilcox, Jeffrey**
El Dorado Hills, CA 95762 (US)
- **Derr, Michael**
El Dorado Hills, CA 95762 (US)
- **Songer, Neil**
Santa Clara, CA 95051 (US)
- **Forbell, Craig**
Los Gatos, CA 95032 (US)

(30) Priority: **26.06.2013 US 201313927746**

(71) Applicant: **Intel Corporation**
Santa Clara, CA 95054 (US)

(74) Representative: **Hufton, David Alan**
HGF Limited
Fountain Precinct
Balm Green
Sheffield S1 2JA (GB)

(72) Inventors:
• **Cooper, Barnes**
Tigard, OR 97223 (US)

(54) **Controlling reduced power states using platform latency tolerance**

(57) In an embodiment, a processor includes a plurality of cores and power management logic. The power management logic may be to, in response to a first break event during a reduced power state in the processor, set an exit timer based on a platform latency tolerance, block

a first plurality of break events from interrupting the reduced power state, and in response to a expiration of the exit timer, terminate the reduced power state. Other embodiments are described and claimed.




FIG. 1

Description

Technical Field

[0001] Embodiments relate generally to power management of electronic devices.

Background

[0002] Conventionally, an electronic device may include one or more power states. Each power state may correspond to a particular performance level and power consumption. Further, each power state may be associated with a particular level of power consumption. The use of such power states may decrease the total amount of electrical power consumed by the electronic device.

Brief Description of the Drawings

[0003]

FIG. 1 is a block diagram in accordance with one or more embodiments.

FIGs. 2A-2B are timing diagrams in accordance with one or more embodiments.

FIG. 3 is a sequence in accordance with one or more embodiments.

FIG. 4 is a block diagram of a processor in accordance with an embodiment of the present invention.

FIG. 5 is a block diagram of a multi-domain processor in accordance with another embodiment of the present invention.

FIG. 6 is a block diagram of an embodiment of a processor including multiple cores.

FIG. 7 is a block diagram of a system in accordance with an embodiment of the present invention.

FIG. 8 is a block diagram of a system on a chip in accordance with an embodiment of the present invention.

Detailed Description

[0004] Some computing systems include functionality to determine a power state based on latency tolerance. Conventionally, such systems may select the deepest (i.e., least power consumption) power state to be entered by determining an overall system latency tolerance, and then rounding off to the deepest power state that has a response time less than the overall platform latency tolerance. Further, such systems typically return to a full power state upon receiving a break event. Thus, conven-

tional systems may not use all available time in a low power state. Therefore, some potential power savings may not be realized.

[0005] In accordance with some embodiments, a computer system may include functionality to defer an exit from a reduced power state based on a negotiated platform latency tolerance. Further, the computer system may include functionality to block some types of break events during the period of the reduced power state. In this manner, a forced delay may be imposed on some break events. Therefore, embodiments may increase the time spent in the reduced power state, and thereby reduce overall power consumption.

[0006] Although the following embodiments are described with reference to energy conservation and energy efficiency in specific integrated circuits, such as in computing platforms or processors, other embodiments are applicable to other types of integrated circuits and logic devices. Similar techniques and teachings of embodiments described herein may be applied to other types of circuits or semiconductor devices that may also benefit from better energy efficiency and energy conservation. For example, the disclosed embodiments are not limited to any particular type of computer systems, and may be also used in other devices, such as handheld devices, systems on chip (SoCs), and embedded applications. Some examples of handheld devices include cellular phones, Internet protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Embedded applications typically include a microcontroller, a digital signal processor (DSP), network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that can perform the functions and operations taught below.

[0007] Moreover, the apparatus, methods, and systems described herein are not limited to physical computing devices, but may also relate to software optimizations for energy conservation and efficiency. As will become readily apparent in the description below, the embodiments of methods, apparatus, and systems described herein (whether in reference to hardware, firmware, software, or a combination thereof) are vital to a 'green technology' future, such as for power conservation and energy efficiency in products that encompass a large portion of the US economy.

[0008] Note that embodiments described herein may be independent of and/or complementary to an operating system (OS)-based mechanism, such as the Advanced Configuration and Power Interface (ACPI) standard (e.g., Rev. 3.0b, published October 10, 2006). According to ACPI, a processor can operate at various performance states or levels, namely from P0 to PN. In general, the P1 performance state may correspond to the highest guaranteed performance state that can be requested by an OS. In addition to this P1 state, the OS can further request a higher performance state, namely a P0 state. This P0 state may thus be an opportunistic state in which, when power and/or thermal budget is available, proces-

sor hardware can configure the processor or at least portions thereof to operate at a higher than guaranteed frequency. In many implementations a processor can include multiple so-called bin frequencies above a guaranteed maximum frequency, also referred to as a P1 frequency. In addition, according to ACPI, a processor can operate at various power states or levels. With regard to power states, ACPI specifies different power consumption states, generally referred to as C-states, C0, C1 to Cn states. When a core is active, it runs at a C0 state, and when the core is idle it may be placed in a core low power state, also called a core non-zero C-state (e.g., C1-C6 states), with each C-state being at a lower power consumption level (such that C6 is a deeper low power state than C1, and so forth).

[0009] Referring to FIG. 1, shown is a block diagram of a system 100 in accordance with one or more embodiments. In some embodiments, the system 100 may be all or a portion of an electronic device or component. For example, the system 100 may be a cellular telephone, a computer, a server, a network device, a controller, an appliance, etc. In another example, the system 100 may be a multicore processor or a System on a Chip (SoC).

[0010] As shown in FIG. 1, the system 100 may include processor(s) 110, PM logic 120, memory 150, chipset 160, and devices 130a-130n. The memory 150 may be any type of computer memory (e.g., dynamic random access memory (DRAM), static random-access memory (SRAM), non-volatile memory, etc.). In some embodiments, the processor(s) 110 may include multiple cores. Further, in embodiments in which the system 100 is a processor or SoC, the processor(s) 110 may be processing cores.

[0011] In some embodiments, the devices 130a-130n may be any hardware/software components associated with the system 100. The devices 130a-130n may include external devices coupled to the system 100, internal devices installed in the system 100, software installed or executing on the system 100, etc. For example, in some embodiments, the devices 130a-130n may include one or more of a peripheral device, a printer, a scanner, a storage drive, a camera, a network adapter, a host controller, a memory controller, a network controller, a graphics controller, a hard disk controller (HDD), an audio controller, a software application, a device driver, an operating system, etc.

[0012] In one or more embodiments, the chipset 160 may include functionality to support the processor 110, memory 150, and/or devices 130a-130n. For example, the chipset 160 may include functionality such as input/output control, memory access, display/audio interface, clocking, etc. In some embodiments, the chipset 160 may include a platform controller hub (PCH).

[0013] In one or more embodiments, the PM logic 120 may include functionality to receive latency time information from the various components of the system 100 (e.g., devices 130, memory 150, chipset 160, etc.). In some embodiments, the latency time information for each com-

ponent may be based at least in part on the maximum response latency that the component may tolerate without adversely affecting its functionality and/or performance. Further, in some embodiments, the PM logic 120 may receive the latency time information via a latency tolerance messaging (LTM) system (e.g., using specialized notification packets to communicate latency tolerance information).

[0014] In one or more embodiments, the PM logic 120 may include functionality to determine an overall platform latency tolerance (PLT) based on the received latency time information. For example, in some embodiments, the PM logic 120 may compare the latency time requirements of various components of the system 100, and may determine the PLT based on the tightest latency constraint (e.g., the component having the smallest latency time).

[0015] In one or more embodiments, the PM logic 120 may control the duration of a power state based on the negotiated PLT. For example, the PM logic 120 may set an exit timer to transition out of a reduced power state in the system 100. In some embodiments, the PM logic 120 may initiate the exit timer in response to break event. Further, in some embodiments, the PM logic 120 may set the exit timer to a time value equal to the PLT minus a wake time (e.g., the time required to return to a normal power state). Thus, in some embodiments, the PM logic 120 may enable the reduced power state to be maintained as long as possible under the constraint of the PLT.

[0016] In one or more embodiments, the PM logic 120 may include functionality to block break events. For example, in some embodiments, the PM logic 120 may block a first type of break event from interrupting a current power state until the exit timer expires. Further, in some embodiments, the PM logic 120 may allow a second type of break event to interrupt or terminate the current power state before the exit timer expires. In some embodiments, the first type of break event may be non-critical or deferrable break events. For example, the first type of break event may include direct memory access (DMA) accesses, non-critical interrupts, non-critical device traffic, etc. Further, in some embodiments, the second type of break event may be critical or non-deferrable break events. For example, the second type of break event may include critical timers (e.g., advanced programmable interrupt controller (APIC) timer, time stamp counter (TSC) deadline timer, virtualization timers, high precision event timer (HPET)), critical interrupts (e.g., vertical blanking interval interrupt (VBI)), critical device traffic, etc. In one or more embodiments, the PM logic 120 may unblock the first type of break event after terminating or exiting a reduced power state. The unblocked break events may then be granted or processed in a normal manner.

[0017] Note that, while the PM logic 120 is depicted in FIG. 1 as being separate from other components of the system 100, embodiments are not limited in this regard. For example, in some embodiments, all or a part of the PM logic 120 may be included in the processor 110 and/or

the chipset 160. The functionality of the PM logic 120 is described further below with reference to FIGs. 2A, 2B, and 3.

[0018] Referring now to FIG. 2A, shown are example timing diagrams in accordance with one or more embodiments. In particular, FIG. 2A illustrates an example of the functionality of the PM logic 120 shown in FIG. 1.

[0019] As shown, FIG. 2A includes a power diagram 220, a request diagram 230, and a grant diagram 240, which all correspond to the same time period. The power diagram 220 illustrates the power state (PS) of a device (along the vertical axis) at various points in time (along the horizontal axis). Initially, the device (e.g., system 100 shown in FIG. 1) is at a normal power state PS0. As shown, a transition to a reduced power state PS1 is completed at time T0. The device remains at the reduced power state PS1 between times T0 and T3. Further, a transition to the normal power state PS0 is initiated at time T3, and is completed at time T4. The device remains at the normal power state PS0 until time T5, at which time a transition back to the reduced power state PS1 is initiated.

[0020] In the example of FIG. 2A, the request diagram 230 shows non-critical break events 201-205 at the times that they are requested. Further, the grant diagram 240 shows the same break events 201-205 at the times that they are granted (i.e., processed or executed).

[0021] As shown, the non-critical break event 201 is requested prior to T1, and thus is requested during the normal power state PS0. Further, the non-critical break event 206 is requested between T4 and T5, and thus is also requested during the normal power state PS0. In some embodiments, during a normal power state, the PM logic 120 does not block or defer non-critical break events. Thus, as shown in the grant diagram 240, the non-critical break events 201 and 206 are not deferred, and are thus granted at substantially the same times that they were requested.

[0022] As shown, the non-critical break event 202 is requested at time T1, and is the first non-critical event to occur during the reduced power state PS1. In one or more embodiments, the PM logic 120 may respond to the first non-critical event to occur during a reduced power state by scheduling an exit from the reduced power state based on a platform latency time ("PLT") for the device. For example, in some embodiments, the PM logic 120 may set an exit timer equal to the PLT minus a response time ("RT1") to transition from PS1 to PS0. Thus, in the example shown in FIG. 2A, the transition from PS1 to PS0 is initiated at time T3, and is completed at time T4 (i.e., after the response time RT1).

[0023] As shown, in this example, the non-critical break events 202, 203, 204 and 205 are requested during the reduced power state PS1 (i.e., between T0 and T3). In some embodiments, the PM logic 120 may cause non-critical break events to be deferred until transitioning out of a reduced power state. Thus, as shown in the grant diagram 240, the non-critical break events 202, 203, 204

and 205 are deferred until time T4 (i.e., when the device fully returns to the normal power state PS0). In this manner, a forced delay equal is imposed on the non-critical break events 202, 203, 204 and 205.

[0024] Note that, conventionally, the transition from PS1 to PS0 may be initiated at the request time of the first non-critical event (e.g., the request time T1 for the non-critical event 202), and may thus be completed at time T2. Thus, by deferring the transition to complete at T4 rather than at T2, the PM logic 120 may enable the reduced power state PS1 to be maintained for an additional time ("AT"). In some embodiments, the additional time AT in the reduced power state PS1 may result in reduced power consumption for the device.

[0025] Referring now to FIG. 2B, shown are example timing diagrams in accordance with one or more embodiments. In particular, FIG. 2B illustrates an example similar to the example shown in FIG. 2A. For instance, FIG. 2B includes a power diagram 225, a request diagram 235, and a grant diagram 245, corresponding respectively to diagrams 220, 230, and 240 of FIG. 2A.

[0026] Assume that, in the example shown in FIG. 2B, the reduced power state PS1 is again initiated at time T0. Further, as in the example shown in FIG. 2A, the non-critical break event 202 is the first non-critical event to occur during the reduced power state PS1. Thus, an exit timer is again set to complete the transition by the PLT (i.e., at time T4). However, as shown in the request diagram 235, the critical break event 207 is requested at time T6. As discussed above, in some embodiments, the PM logic 120 may terminate a reduced power state in response to a critical break event. Thus, as shown in the power diagram 225, a transition out of PS1 is initiated at T6, and is then completed at time T7. Further, as shown in the grant diagram 245, the critical break event 207 and the non-critical break events 202, 203, 204, and 205 are granted at time T7.

[0027] Referring now to FIG. 3, shown is a sequence 300 for managing a power state, in accordance with one or more embodiments. In one or more embodiments, the sequence 300 may be part of the PM logic 120 shown in FIG. 1. The sequence 300 may be implemented in hardware, software, and/or firmware. In firmware and software embodiments it may be implemented by computer executed instructions stored in a non-transitory computer readable medium, such as an optical, semiconductor, or magnetic storage device.

[0028] At step 310, a reduced power state may be initiated. For example, referring to FIGs. 1 and 2A, the PM logic 120 may initiate a transition from the normal power state PS0 to the reduced power state PS1 at time T1. In some embodiments, the reduced power state may associated with a lower power consumption level than a normal or higher power state.

[0029] At step 320, a determination is made about whether the reduced power state is below a defined threshold level. For example, referring to FIG. 1, the PM logic 120 may determine whether the reduced power

state is a deeper (e.g., provides less power consumption) than a specific power state. In some embodiments, the defined threshold level may correspond to, e.g., the C2 power state. Further, in some embodiments, the threshold level may be defined based on the amount of power efficiency which is estimated to be available in a specific reduced power state.

[0030] If it is determined at step 320 that the new power state is not below the defined threshold level, then the sequence 300 may be terminated. However, if it is determined at step 320 that the new power state is below the defined threshold level, then the sequence 300 continues at step 325.

[0031] At step 325, a determination is made about whether bus traffic is below a defined threshold level. For example, referring to FIG. 1, the PM logic 120 may determine whether a bus and/or backbone of the system 100 has not had any traffic for at least a minimum time period (e.g., 5 microseconds, 10 microseconds, etc.).

[0032] If it is determined at step 325 that the bus traffic is not below the defined threshold level, then the sequence 300 may terminate. However, if it is determined at step 325 that the bus traffic is below the defined threshold level, then the sequence 300 continues at step 327.

[0033] At step 327, non-critical break events may be blocked. For example, referring to FIG. 1, the PM logic 120 may block non-critical break events from interrupting the reduced power state. Such blocking of non-critical break events may include blocking data in a bus or interface (e.g., the internal bus of the I/O controller or backbone) of the system 100. Further, the blocked break events may include, e.g., a DMA transfer, non-critical interrupts, non-critical device traffic, etc.

[0034] At step 330, a determination is made about whether a non-critical break event has occurred during the reduced power state (initiated at step 310). For example, referring to FIGs. 1 and 2A, the PM logic 120 may detect the first non-critical break event 202 to occur during the reduced power state PS1. If it is determined at step 330 that a non-critical break event has not occurred during the reduced power state, step 330 may be repeated to continue monitoring for a non-critical break event. However, if it is determined at step 330 that a non-critical break event has occurred, then the sequence 300 continues at step 335.

[0035] At step 335, a determination is made about whether an exit timer would expire prior to any existing timer. In some embodiments, the exit timer may be based on a PLT value. For example, referring to FIG. 1, the PM logic 120 may determine the PLT for the system 100 (e.g., based on the component having the smallest latency time). The PM logic 120 may calculate the period of the exit timer as the PLT value minus a wake time. Further, the PM logic 120 may compare the exit timer to the existing timer. If the exit timer is shorter than the existing timer, the PM logic 120 may determine that the exit timer would expire sooner than the existing timer.

[0036] If it is determined at step 335 that the exit timer

would expire prior to any existing timer, then at step 340, the exit timer may be initiated. For example, referring to FIG. 1, the PM logic 120 may set an exit timer equal to the PLT value minus a wake time.

[0037] At step 345, a determination is made about whether the exit timer has expired. For example, referring to FIG. 1, the PM logic 120 may determine whether the exit timer has expired. If it is determined at step 345 that the exit timer has expired, then the sequence 300 continues at step 370 (described below). Otherwise, if it is determined at step 345 that the exit timer has not expired, then the sequence 300 continues at step 350.

[0038] At step 350, a determination is made about whether a critical break event has occurred. For example, referring to FIGs. 1 and 2B, the PM logic 120 may determine whether any critical break events (e.g., critical break event 207) have occurred. If it is determined at step 350 that a critical break event has not occurred, the sequence 300 may return to step 345 to continue monitoring expiration of the exit timer. However, if it is determined at step 350 that a critical break event has occurred, then the sequence 300 continues at step 370 (described below).

[0039] Returning to step 335, if it is determined that the exit timer would not expire prior to an existing timer, then at step 360, a determination is made about whether the existing timer has expired. If it is determined at step 360 that the existing timer has expired, then the sequence 300 continues at step 370 (described below). Otherwise, if it is determined at step 360 that the existing timer has not expired, then the sequence 300 continues at step 365.

[0040] At step 365, a determination is made about whether a critical break event has occurred. If it is determined at step 365 that a critical break event has not occurred, the sequence 300 may return to step 360 to continue monitoring expiration of the existing timer. However, if it is determined at step 365 that a critical break event has occurred, then the sequence 300 continues at step 370.

[0041] At step 370, the reduced power state (initiated at step 310) may be terminated. For example, referring to FIGs. 1 and 2A, the PM logic 120 may initiate a transition from the reduced power state PS1 to the normal power state PS0 at time T3.

[0042] At step 380, non-critical events may be unblocked. For example, referring to FIG. 1, the PM logic 120 may unblock non-critical break events. Further, in some embodiments, any deferred events may be handled/processed. For example, referring to FIG. 2A, the deferred break events 202, 203, 204, and 205 may be granted or processed when the transition to the normal power state PS0 is completed at time T4. After step 380, the sequence 300 may terminate.

[0043] Note that the examples shown in FIGs. 1, 2A, 2B, and 3 are provided for the sake of illustration, and are not intended to limit any embodiments. For instance, while embodiments may be shown in simplified form for the sake of clarity, embodiments may include any number

and/or arrangement of additional components (e.g., processors, cores, buses, storage media, connectors, power components, buffers, interfaces, etc.). Further, in some embodiments, the system 100 may be a multi-core processor or a System on a Chip (SoC) integrated on a single die or integrated circuit. It is contemplated that specifics in the examples shown in FIGs. 1, 2A, 2B, and 3 may be used anywhere in one or more embodiments.

[0044] Referring now to FIG. 4, shown is a block diagram of a processor in accordance with an embodiment of the present invention. As shown in FIG. 4, the processor 400 may be a multicore processor including first die 405 having a plurality of cores 410a - 410n of a core domain. The various cores 410a - 410n may be coupled via an interconnect 415 to a system agent or uncore domain that includes various components. As seen, the uncore domain may include a shared cache 430. In addition, the uncore may include an integrated memory controller 440, a power control unit (PCU) 470, and various interfaces 450.

[0045] With further reference to FIG. 4, the processor 400 may communicate with a system memory 460, e.g., via a memory bus. In addition, by interfaces 450, connection can be made to another processor, or various off-package components such as peripheral devices, mass storage and so forth. In some embodiments, the processor 400 may include some or all of the functionality of the PM logic 120 shown in FIG. 1. While shown with this particular implementation in the embodiment of FIG. 4, the scope of the present invention is not limited in this regard.

[0046] Referring now to FIG. 5, shown is a block diagram of a multi-domain processor in accordance with another embodiment of the present invention. As shown in the embodiment of FIG. 5, processor 500 includes multiple domains. Specifically, a core domain 510 can include a plurality of cores 510a-510n, a graphics domain 520 can include one or more graphics engines, and a system agent domain 550 may further be present. Note that while only shown with three domains, understand the scope of the present invention is not limited in this regard and additional domains can be present in other embodiments. For example, multiple core domains may be present each including at least one core.

[0047] In general, each core 510 may further include low level caches in addition to various execution units and additional processing elements. In turn, the various cores may be coupled to each other and to a shared cache memory formed of a plurality of units of a last level cache (LLC) 540a - 540n. In various embodiments, LLC 540 may be shared amongst the cores and the graphics engine, as well as various media processing circuitry.

[0048] As seen, a ring interconnect 530 thus couples the cores together, and provides interconnection between the cores, graphics domain 520 and system agent circuitry 550. In some embodiments, the ring interconnect 530 may be a multiplexor or crossbar device. In the embodiment of FIG. 5, system agent domain 550 may in-

clude display controller 552 which may provide control of and an interface to an associated display. As further seen, system agent domain 550 may also include a power control unit 555 to allocate power to the CPU and non-CPU domains.

[0049] As further seen in FIG. 5, processor 500 can further include an integrated memory controller (IMC) 570 that can provide for an interface to a system memory, such as a dynamic random access memory (DRAM). Multiple interfaces 580a - 580n may be present to enable interconnection between the processor and other circuitry. For example, in one embodiment at least one direct media interface (DMI) interface may be provided as well as one or more Peripheral Component Interconnect Express (PCI Express™ (PCIe™)) interfaces. Still further, to provide for communications between other agents such as additional processors or other circuitry, one or more interfaces in accordance with an Intel® Quick Path Interconnect (QPI) protocol may also be provided. As further seen, a peripheral controller hub (PCH) 590 may also be present within the processor 500, and can be implemented on a separate die, in some embodiments. Alternatively, in some embodiments, the PCH 590 may be external to the processor 500. In some embodiments, the processor 500 may include some or all of the functionality of the PM logic 120 shown in FIG. 1. Although shown at this high level in the embodiment of FIG. 5, understand the scope of the present invention is not limited in this regard.

[0050] Referring to FIG. 6, an embodiment of a processor including multiple cores is illustrated. Processor 1100 includes any processor or processing device, such as a microprocessor, an embedded processor, a digital signal processor (DSP), a network processor, a handheld processor, an application processor, a co-processor, a system on a chip (SOC), or other device to execute code. Processor 1100, in one embodiment, includes at least two cores-cores 1101 and 1102, which may include asymmetric cores or symmetric cores (the illustrated embodiment). However, processor 1100 may include any number of processing elements that may be symmetric or asymmetric. In some embodiments, the processor 1100 may include some or all of the functionality of the PM logic 120 shown in FIG. 1.

[0051] In one embodiment, a processing element refers to hardware or logic to support a software thread. Examples of hardware processing elements include: a thread unit, a thread slot, a thread, a process unit, a context, a context unit, a logical processor, a hardware thread, a core, and/or any other element, which is capable of holding a state for a processor, such as an execution state or architectural state. In other words, a processing element, in one embodiment, refers to any hardware capable of being independently associated with code, such as a software thread, operating system, application, or other code. A physical processor typically refers to an integrated circuit, which potentially includes any number of other processing elements, such as cores or hardware

threads.

[0052] A core often refers to logic located on an integrated circuit capable of maintaining an independent architectural state, wherein each independently maintained architectural state is associated with at least some dedicated execution resources. In contrast to cores, a hardware thread typically refers to any logic located on an integrated circuit capable of maintaining an independent architectural state, wherein the independently maintained architectural states share access to execution resources. As can be seen, when certain resources are shared and others are dedicated to an architectural state, the line between the nomenclature of a hardware thread and core overlaps. Yet often, a core and a hardware thread are viewed by an operating system as individual logical processors, where the operating system is able to individually schedule operations on each logical processor.

[0053] Physical processor 1100, as illustrated in FIG. 6, includes two cores, cores 1101 and 1102. Here, cores 1101 and 1102 are considered symmetric cores, i.e. cores with the same configurations, functional units, and/or logic. In another embodiment, core 1101 includes an out-of-order processor core, while core 1102 includes an in-order processor core. However, cores 1101 and 1102 may be individually selected from any type of core, such as a native core, a software managed core, a core adapted to execute a native instruction set architecture (ISA), a core adapted to execute a translated ISA, a co-designed core, or other known core. Yet to further the discussion, the functional units illustrated in core 1101 are described in further detail below, as the units in core 1102 operate in a similar manner.

[0054] As shown, core 1101 includes two hardware threads 1101 a and 1101 b, which may also be referred to as hardware thread slots 1101 a and 1101 b. Therefore, software entities, such as an operating system, in one embodiment potentially view processor 1100 as four separate processors, i.e., four logical processors or processing elements capable of executing four software threads concurrently. As alluded to above, a first thread is associated with architecture state registers 1101 a, a second thread is associated with architecture state registers 1101 b, a third thread may be associated with architecture state registers 1102a, and a fourth thread may be associated with architecture state registers 1102b. Here, each of the architecture state registers (1101 a, 1101 b, 1102a, and 1102b) may be referred to as processing elements, thread slots, or thread units, as described above.

[0055] As illustrated, architecture state registers 1101 a are replicated in architecture state registers 1101 b, so individual architecture states/contexts are capable of being stored for logical processor 1101 a and logical processor 1101 b. In core 1101, other smaller resources, such as instruction pointers and renaming logic in allocator and renamer block 1130 may also be replicated for threads 1101 a and 1101 b. Some resources, such as

re-order buffers in reorder/retirement unit 1135, ILTB 1120, load/store buffers, and queues may be shared through partitioning. Other resources, such as general purpose internal registers, page-table base register(s), low-level data-cache and data-TLB 1115, execution unit(s) 1140, and portions of out-of-order unit 1135 are potentially fully shared.

[0056] Processor 1100 often includes other resources, which may be fully shared, shared through partitioning, or dedicated by/to processing elements. In FIG. 6, an embodiment of a purely exemplary processor with illustrative logical units/resources of a processor is illustrated. Note that a processor may include, or omit, any of these functional units, as well as include any other known functional units, logic, or firmware not depicted. As illustrated, core 1101 includes a simplified, representative out-of-order (OOO) processor core. But an in-order processor may be utilized in different embodiments. The OOO core includes a branch target buffer 1120 to predict branches to be executed/taken and an instruction-translation buffer (I-TLB) 1120 to store address translation entries for instructions.

[0057] Core 1101 further includes decode module 1125 coupled to fetch unit 1120 to decode fetched elements. Fetch logic, in one embodiment, includes individual sequencers associated with thread slots 1101 a, 1101 b, respectively. Usually core 1101 is associated with a first ISA, which defines/specifies instructions executable on processor 1100. Often machine code instructions that are part of the first ISA include a portion of the instruction (referred to as an opcode), which references/specifies an instruction or operation to be performed. Decode logic 1125 includes circuitry that recognizes these instructions from their opcodes and passes the decoded instructions on in the pipeline for processing as defined by the first ISA. As a result of the recognition by decoders 1125, the architecture or core 1101 takes specific, predefined actions to perform tasks associated with the appropriate instruction. It is important to note that any of the tasks, blocks, operations, and methods described herein may be performed in response to a single or multiple instructions; some of which may be new or old instructions.

[0058] In one example, allocator and renamer block 1130 includes an allocator to reserve resources, such as register files to store instruction processing results. However, threads 1101 a and 1101 b are potentially capable of out-of-order execution, where allocator and renamer block 1130 also reserves other resources, such as reorder buffers to track instruction results. Unit 1130 may also include a register renamer to rename program/instruction reference registers to other registers internal to processor 1100. Reorder/retirement unit 1135 includes components, such as the reorder buffers mentioned above, load buffers, and store buffers, to support out-of-order execution and later in-order retirement of instructions executed out-of-order.

[0059] Scheduler and execution unit(s) block 1140, in one embodiment, includes a scheduler unit to schedule

instructions/operation on execution units. For example, a floating point instruction is scheduled on a port of an execution unit that has an available floating point execution unit. Register files associated with the execution units are also included to store information instruction processing results. Exemplary execution units include a floating point execution unit, an integer execution unit, a jump execution unit, a load execution unit, a store execution unit, and other known execution units.

[0060] Lower level data cache and data translation buffer (D-TLB) 1150 are coupled to execution unit(s) 1140. The data cache is to store recently used/operated on elements, such as data operands, which are potentially held in memory coherency states. The D-TLB is to store recent virtual/linear to physical address translations. As a specific example, a processor may include a page table structure to break physical memory into a plurality of virtual pages.

[0061] Here, cores 1101 and 1102 share access to higher-level or further-out cache 1110, which is to cache recently fetched elements. Note that higher-level or further-out refers to cache levels increasing or getting further away from the execution unit(s). In one embodiment, higher-level cache 1110 is a last-level data cache-last cache in the memory hierarchy on processor 1100—such as a second or third level data cache. However, higher level cache 1110 is not so limited, as it may be associated with or includes an instruction cache. A trace cache—a type of instruction cache—instead may be coupled after decoder 1125 to store recently decoded traces. In the depicted configuration, processor 1100 also includes bus interface module 1105 and a power controller 1160, which may perform power management in accordance with an embodiment of the present invention.

[0062] Historically, controller 1170 has been included in a computing system external to processor 1100. In this scenario, bus interface 1105 is to communicate with devices external to processor 1100, such as system memory 1175, a chipset (often including a memory controller hub to connect to memory 1175 and an I/O controller hub to connect peripheral devices), a memory controller hub, a northbridge, or other integrated circuit. And in this scenario, bus 1105 may include any known interconnect, such as multi-drop bus, a point-to-point interconnect, a serial interconnect, a parallel bus, a coherent (e.g. cache coherent) bus, a layered protocol architecture, a differential bus, and a GTL bus.

[0063] Memory 1175 may be dedicated to processor 1100 or shared with other devices in a system. Common examples of types of memory 1175 include DRAM, SRAM, non-volatile memory (NV memory), and other known storage devices. Note that device 1180 may include a graphic accelerator, processor or card coupled to a memory controller hub, data storage coupled to an I/O controller hub, a wireless transceiver, a flash device, an audio controller, a network controller, or other known device.

[0064] Note however, that in the depicted embodiment,

the controller 1170 is illustrated as part of processor 1100. Recently, as more logic and devices are being integrated on a single die, such as SOC, each of these devices may be incorporated on processor 1100. For example in one embodiment, memory controller hub 1170 is on the same package and/or die with processor 1100. Here, a portion of the core (an on-core portion) includes one or more controller(s) 1170 for interfacing with other devices such as memory 1175 or a graphics device 1180. The configuration including an interconnect and controllers for interfacing with such devices is often referred to as an on-core (or un-core configuration). As an example, bus interface 1105 includes a ring interconnect with a memory controller for interfacing with memory 1175 and a graphics controller for interfacing with graphics processor 1180. Yet, in the SOC environment, even more devices, such as the network interface, coprocessors, memory 1175, graphics processor 1180, and any other known computer devices/interface may be integrated on a single die or integrated circuit to provide small form factor with high functionality and low power consumption.

[0065] Embodiments may be implemented in many different system types. Referring now to FIG. 7, shown is a block diagram of a system in accordance with an embodiment of the present invention. As shown in FIG. 7, multiprocessor system 600 is a point-to-point interconnect system, and includes a first processor 670 and a second processor 680 coupled via a point-to-point interconnect 650. As shown in FIG. 7, each of processors 670 and 680 may be multicore processors, including first and second processor cores (i.e., processor cores 674a and 674b and processor cores 684a and 684b), although potentially many more cores may be present in the processors. In some embodiments, the processors 670, 680 may include some or all of the functionality of the PM logic 120 shown in FIG. 1.

[0066] Still referring to FIG. 7, first processor 670 further includes a memory controller hub (MCH) 672 and point-to-point (P-P) interfaces 676 and 678. Similarly, second processor 680 includes a MCH 682 and P-P interfaces 686 and 688. As shown in FIG. 7, MCH's 672 and 682 couple the processors to respective memories, namely a memory 632 and a memory 634, which may be portions of system memory (e.g., DRAM) locally attached to the respective processors. First processor 670 and second processor 680 may be coupled to a chipset 690 via P-P interconnects 652 and 654, respectively. As shown in FIG. 7, chipset 690 includes P-P interfaces 694 and 698.

[0067] Furthermore, chipset 690 includes an interface 692 to couple chipset 690 with a high performance graphics engine 638, by a P-P interconnect 639. In turn, chipset 690 may be coupled to a first bus 616 via an interface 696. As shown in FIG. 7, various input/output (I/O) devices 614 may be coupled to first bus 616, along with a bus bridge 618 which couples first bus 616 to a second bus 620. Various devices may be coupled to second bus 620 including, for example, a keyboard/mouse 622, com-

munication devices 626 and a data storage unit 628 such as a disk drive or other mass storage device which may include code 630, in one embodiment. Further, an audio I/O 624 may be coupled to second bus 620. Embodiments can be incorporated into other types of systems including mobile devices such as a smart cellular telephone, tablet computer, netbook, Ultrabook™, or so forth.

[0068] It should be understood that a processor core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

[0069] Any processor described herein may be a general-purpose processor, such as a Core™ i3, i5, i7, 2 Duo and Quad, Xeon™, Itanium™, XScale™ or StrongARM™ processor, which are available from Intel Corporation, of Santa Clara, Calif. Alternatively, the processor may be from another company, such as ARM Holdings, Ltd, MIPS, etc.. The processor may be a special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, co-processor, embedded processor, or the like. The processor may be implemented on one or more chips. The processor may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

[0070] It is contemplated that the processors described herein are not limited to any system or device. Other system designs and configurations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs, switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, are also suitable. In general, a huge variety of systems or electronic devices capable of incorporating a processor and/or other execution logic as disclosed herein are generally suitable.

[0071] Turning next to FIG. 8, an embodiment of a system on-chip (SOC) design in accordance with the inventions is depicted. As a specific illustrative example, SOC 2000 is included in user equipment (UE). In one embodiment, UE refers to any device to be used by an end-user to communicate, such as a hand-held phone, smart-phone, tablet, ultra-thin notebook, notebook with broadband adapter, or any other similar communication device. Often a UE connects to a base station or node, which potentially corresponds in nature to a mobile station (MS) in a GSM network. In some embodiments, SOC 2000 may include some or all of the functionality of the PM

logic 120 shown in FIG. 1.

[0072] Here, SOC 2000 includes 2 cores-2006 and 2007. The cores 2006 and 2007 may conform to an Instruction Set Architecture, such as an Intel® Architecture Core™-based processor, an Advanced Micro Devices, Inc. (AMD) processor, a MIPS-based processor, an ARM-based processor design, or a customer thereof, as well as their licensees or adopters. Cores 2006 and 2007 are coupled to cache control 2008 that is associated with bus interface unit 2009 and L2 cache 2011 to communicate with other parts of system 2000. Interconnect 2010 includes an on-chip interconnect, such as an IOSF, AMBA, or any other interconnect, which potentially implements one or more aspects of the described invention.

[0073] Interface 2010 provides communication channels to the other components, such as a Subscriber Identity Module (SIM) 2030 to interface with a SIM card, a boot rom 2035 to hold boot code for execution by cores 2006 and 2007 to initialize and boot SOC 2000, a SDRAM controller 2040 to interface with external memory (e.g. DRAM 2060), a flash controller 2045 to interface with non-volatile memory (e.g. Flash 2065), a peripheral control Q1650 (e.g. Serial Peripheral Interface) to interface with peripherals, video codecs 2020 and Video interface 2025 to display and receive input (e.g. touch enabled input), GPU 2015 to perform graphics related computations, etc.

[0074] In addition, the system illustrates peripherals for communication, such as a Bluetooth module 2070, 3G modem 2075, GPS 2085, and WiFi 2085. Note that a UE includes a radio for communication. As a result, these peripheral communication modules are not all required. However, in a UE some form a radio for external communication is to be included.

[0075] Embodiments may be implemented in code and may be stored on a non-transitory storage medium having stored thereon instructions which can be used to program a system to perform the instructions. The storage medium may include, but is not limited to, any type of disk including floppy disks, optical disks, solid state drives (SSDs), compact disk read-only memories (CD-ROMs), compact disk rewritables (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

[0076] The following clauses and/or examples pertain to further embodiments. In one example embodiment may be a processor including a plurality of cores and power management logic. The power management logic may be to: in response to a first break event during a reduced power state in the processor, set an exit timer based on a platform latency tolerance; block a first plu-

rality of break events from interrupting the reduced power state; and in response to a expiration of the exit timer, terminate the reduced power state.

[0077] In an example, the power management logic may be further to: obtain latency tolerance requirements for each of the plurality of hardware devices; and determine the platform latency tolerance using the latency tolerance requirements.

[0078] In an example, the power management logic may be to obtain the latency tolerance requirements using a latency tolerance messaging (LTM) system.

[0079] In an example, the power management logic may be to set the exit timer to a time value based at least in part on the platform latency tolerance and a wake time.

[0080] In an example, the power management logic may be further to unblock the first plurality of break events after termination of the reduced power state.

[0081] In an example, the first plurality of break events comprises non-critical break events. The non-critical break events may include at least one of a direct memory access (DMA) transfer and a non-critical interrupt.

[0082] In an example, the power management logic may be further to, in response to one or more critical break events, terminate the reduced power state prior to the expiration of the exit timer. The one or more critical break events may include at least one of an advanced programmable interrupt controller (APIC) timer, a time stamp counter (TSC) deadline timer, a virtualization timer, a high precision event timer (HPET), and a vertical blanking interval interrupt (VBI).

[0083] In an example, the power management logic may be further to determine whether the reduced power state is below a specific threshold power state.

[0084] In an example, the power management logic may be further to determine whether bus traffic is below a defined threshold level.

[0085] In another example embodiment may be a processor including a plurality of cores and power management logic. The power management logic may be to: determine a platform latency tolerance for a system comprising a plurality of components; set, based on the platform latency tolerance, an exit timer for a reduced power state; defer one or more non-critical break events until a termination of the reduced power state; and terminate the reduced power state in response to an expiration of the exit timer.

[0086] In an example, the power management logic may be further to initiate the reduced power state in the system.

[0087] In an example, the power management logic may be further to: obtain a plurality of latency tolerance requirements from the plurality of components; and determine the platform latency tolerance based on the plurality of latency tolerance requirements.

[0088] In an example, the power management logic may be to set the exit timer by subtracting a wake time from the platform latency tolerance.

[0089] In an example, the power management logic

may be further to terminate the reduced power state in response to one or more critical break events.

[0090] In another example embodiment may be a method, the method including: obtaining, by power management logic of a computing system, latency tolerance information for a plurality of devices associated with the computing system; determining a platform latency tolerance based on the latency tolerance information; initiating an exit timer to a time period based on the platform latency tolerance; delaying at least one non-critical break event while the computer system is in a reduced power state; and in response to an expiration of the exit timer, terminating the reduced power state.

[0091] In an example, the method may further include calculating the time period by subtracting a wake time from the platform latency tolerance, wherein the wake time is an amount of time to transition from the reduced power state to a normal power state.

[0092] In an example, obtaining the latency tolerance information for the plurality of devices may include receiving a plurality of latency tolerance messages from the plurality of devices.

[0093] In an example, the method may further include, in response to a critical break event: terminating the reduced power state prior to the expiration of the exit timer; processing the critical break event; and processing the delayed at least one non-critical break event.

[0094] In an example, the method may further include initiating the reduced power state in the system.

[0095] In an example, the method may further include determining whether the reduced power state is below a specific threshold power state.

[0096] In an example, the method may further include determining whether bus traffic is below a defined threshold level.

[0097] In an example, the at least one non-critical break event is one of a direct memory access (DMA) transfer and a non-critical interrupt.

[0098] References throughout this specification to "one embodiment" or "an embodiment" mean that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one implementation encompassed within the present invention. Thus, appearances of the phrase "one embodiment" or "in an embodiment" are not necessarily referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be instituted in other suitable forms other than the particular embodiment illustrated and all such forms may be encompassed within the claims of the present application.

[0099] While the present invention has been described with respect to a limited number of embodiments for the sake of illustration, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

Claims**1.** A processor comprising:

a plurality of cores;
power management logic to:

in response to a first break event during a reduced power state in the processor, set an exit timer based on a platform latency tolerance;

block a first plurality of break events from interrupting the reduced power state; and
in response to an expiration of the exit timer, terminate the reduced power state.

2. The processor of claim 1, wherein the power management logic is further to:

obtain latency tolerance requirements for each of the plurality of hardware devices; and
determine the platform latency tolerance using the latency tolerance requirements.

3. The processor of claim 2, wherein the power management logic is to obtain the latency tolerance requirements using a latency tolerance messaging (LTM) system.**4.** The processor of claim 1, wherein the power management logic is to set the exit timer to a time value based at least in part on the platform latency tolerance and a wake time.**5.** The processor of claim 1, wherein the first plurality of break events comprises non-critical break events.**6.** A method, comprising:

obtaining, by power management logic of a computing system, latency tolerance information for a plurality of devices associated with the computing system;

determining a platform latency tolerance based on the latency tolerance information;

initiating an exit timer to a time period based on the platform latency tolerance;

delaying at least one non-critical break event while the computer system is in a reduced power state; and

in response to an expiration of the exit timer, terminating the reduced power state.

7. The method of claim 6, further comprising:

calculating the time period by subtracting a wake time from the platform latency tolerance, wherein the wake time is an amount of time to transition

from the reduced power state to a normal power state.

8. The method of claim 6, wherein obtaining the latency tolerance information for the plurality of devices comprises receiving a plurality of latency tolerance messages from the plurality of devices.**9.** The method of claim 6, further comprising, in response to a critical break event:

terminating the reduced power state prior to the expiration of the exit timer;
processing the critical break event; and
processing the delayed at least one non-critical break event.

10. The method of claim 6, further comprising:

initiating the reduced power state in the system.

11. The method of claim 6, wherein the at least one non-critical break event is one of a direct memory access (DMA) transfer and a non-critical interrupt.**12.** A communication device arranged to perform the method of any one of claims 6 to 11.**13.** At least one machine readable medium comprising a plurality of instructions that in response to being executed on a computing device, cause the computing device to carry out a method according to any one of claims 6 to 11.**14.** An apparatus for processing instructions, configured to perform the method of any one of claims 6 to 11.**15.** An apparatus comprising means for performing the method of any one of claims 6 to 11.




FIG. 1




FIG. 2A




FIG. 2B




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8



EUROPEAN SEARCH REPORT

Application Number
EP 14 17 4593

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 2012/079482 A1 (WANG REN [US] ET AL) 29 March 2012 (2012-03-29) * paragraph [0011] - paragraph [0028]; figures 1-5 * * paragraph [0033] * * paragraph [0004] *	1-15	INV. G06F1/32
A	EP 2 073 097 A2 (INTEL CORP [US]) 24 June 2009 (2009-06-24) * the whole document *	1-15	
A	US 2010/169684 A1 (JEYASEELAN JAYA L [US] ET AL) 1 July 2010 (2010-07-01) * the whole document *	1-15	
A	US 2012/254644 A1 (KNOLLA WILLIAM [US] ET AL) 4 October 2012 (2012-10-04) * the whole document *	1-15	
A	US 2009/249103 A1 (JEYASEELAN JAYA L [US] ET AL) 1 October 2009 (2009-10-01) * paragraph [0020] *	1-15	TECHNICAL FIELDS SEARCHED (IPC)
A	JP 2006 072698 A (MATSUSHITA ELECTRIC IND CO LTD) 16 March 2006 (2006-03-16) * abstract *	1-15	G06F
A	EP 1 471 431 A1 (FUJITSU LTD [JP]) 27 October 2004 (2004-10-27) * the whole document *	1-15	
A	US 2010/169683 A1 (WANG REN [US] ET AL) 1 July 2010 (2010-07-01) * the whole document *	1-15	
A	US 2013/007492 A1 (SOKOL JR JOSEPH [US] ET AL) 3 January 2013 (2013-01-03) * the whole document *	1-15	
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 29 January 2015	Examiner Absalom, Richard
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 03.82 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 14 17 4593

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

29-01-2015

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012079482 A1	29-03-2012	CN 102597912 A	18-07-2012
		DE 112011103194 T5	27-06-2013
		GB 2497025 A	29-05-2013
		JP 5628431 B2	19-11-2014
		JP 2013543613 A	05-12-2013
		KR 20130048786 A	10-05-2013
		TW 201217973 A	01-05-2012
		US 2012079482 A1	29-03-2012
		US 2014101674 A1	10-04-2014
		WO 2012040068 A1	29-03-2012
EP 2073097 A2	24-06-2009	CN 101458558 A	17-06-2009
		EP 2073097 A2	24-06-2009
		US 2009150696 A1	11-06-2009
US 2010169684 A1	01-07-2010	CN 101853065 A	06-10-2010
		DE 102009060269 A1	01-07-2010
		JP 5385124 B2	08-01-2014
		JP 2010165350 A	29-07-2010
		KR 20100080395 A	08-07-2010
		TW 201040729 A	16-11-2010
		US 2010169684 A1	01-07-2010
		US 2014095908 A1	03-04-2014
US 2012254644 A1	04-10-2012	EP 2691831 A2	05-02-2014
		US 2012254644 A1	04-10-2012
		WO 2012134683 A2	04-10-2012
US 2009249103 A1	01-10-2009	CN 101598969 A	09-12-2009
		CN 103645793 A	19-03-2014
		DE 102009015495 A1	24-12-2009
		GB 2458805 A	07-10-2009
		JP 2010009581 A	14-01-2010
		JP 2012212467 A	01-11-2012
		JP 2014206996 A	30-10-2014
		TW 201003377 A	16-01-2010
		US 2009249103 A1	01-10-2009
		US 2012198248 A1	02-08-2012
JP 2006072698 A	16-03-2006	NONE	
EP 1471431 A1	27-10-2004	DE 60223177 T2	31-07-2008
		EP 1471431 A1	27-10-2004
		JP 4219818 B2	04-02-2009
		US 2004236875 A1	25-11-2004
		WO 03063002 A1	31-07-2003

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 14 17 4593

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

29-01-2015

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010169683 A1	01-07-2010	CN 101916137 A	15-12-2010
		JP 5095718 B2	12-12-2012
		JP 2010191945 A	02-09-2010
		TW 201037506 A	16-10-2010
		US 2010169683 A1	01-07-2010
		US 2013311808 A1	21-11-2013
US 2013007492 A1	03-01-2013	NONE	

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Non-patent literature cited in the description

- Rev. 3.0b, 10 October 2006 [0008]